arXiv:1701.03989v1 [cs.NA] 15 Jan 2017

THE ADAPTIVE S-STEP CONJUGATE GRADIENT METHOD

ERIN CARSON*

Abstract. On modern large-scale parallel computers, the performance of Krylov subspace iterative methods is limited by global synchronization. This has inspired the development of s-step (or communication-avoiding) Krylov subspace method variants, in which iterations are computed in blocks of s. This reformulation can reduce the number of global synchronizations per iteration by a factor of $O(s)$, and has been shown to produce speedups in practical settings.

Although the s-step variants are mathematically equivalent to their classical counterparts, they can behave quite differently in finite precision depending on the parameter s. If s is chosen too large, the s-step method can suffer a convergence delay and a decrease in attainable accuracy relative to the classical method. This makes it difficult for a potential user of such methods - the s value that minimizes the time per iteration may not be the best s for minimizing the overall time-to-solution, and further may cause an unacceptable decrease in accuracy.

Towards improving the reliability and usability of s-step Krylov subspace methods, in this work we derive the *adaptive s-step CG method*, a variable s-step CG method where in block k, the parameter s_k is determined automatically such that a user-specified accuracy is attainable. The method for determining s_k is based on a bound on growth of the residual gap within block k, from which we derive a constraint on the condition numbers of the computed $O(s_k)$-dimensional Krylov subspace bases. The computations required for determining the block size s_k can be performed without increasing the number of global synchronizations per block. Our numerical experiments demonstrate that the adaptive s-step CG method is able to attain up to the same accuracy as classical CG while still significantly reducing the total number of global synchronizations.

1. Introduction.

In this paper, we consider the use of Krylov subspace methods for solving large, sparse linear systems $Ax = b$, where $A \in \mathbb{R}^{n\times n}$. We will focus on the conjugate gradient (CG) method [25], which is used when A is symmetric positive definite. Given an initial approximate solution x_1 and corresponding residual $r_1 = b - Ax_1$, the CG method iteratively updates the approximate solution using a sequence of nested Krylov subspaces $\mathcal{K}_1(A, r_1) \subset \mathcal{K}_2(A, r_1) \subset \cdots$, where

$$\mathcal{K}_i(A, r_1) = \text{span}(r_1, Ar_1, \ldots, A^{i-1}r_1)$$

denotes the i-dimension Krylov subspace with matrix A and starting vector v. In iteration i, the updated approximate solution $x_{i+1} \in x_1 + \mathcal{K}_i(A, r_1)$ is chosen by imposing the Galerkin condition $r_{i+1} = b - Ax_{i+1} \perp \mathcal{K}_i$.

Thus each iteration of CG requires a matrix-vector product with A in order to expand the dimension of the Krylov subspace and a number of inner products to perform the orthogonalization. On modern computer architectures, the speed with which the matrix-vector products and inner products can be computed is limited by communication (i.e., the movement of data). This limits the potential speed of individual iterations attainable by an implementation of CG. To perform a sparse matrix-vector product in parallel, each processor must communicate entries of the source vector and/or the destination vector that it owns to neighboring processors. Inner products require a global synchronization, i.e., the computation can not proceed until all processors have finished their local computation and communicated the result to other processors. For large-scale sparse problems on large-scale machines, the cost of synchronization between parallel processors can dominate the run-time (see, e.g., the exascale computing report [19, pp. 28]).

*Courant Institute of Mathematical Sciences, New York University, New York, NY. (erinc@cims.nyu.edu, http://math.nyu.edu/~erinc/).

Research efforts toward removing the performance bottlenecks caused by communication in CG and other Krylov subspace methods have produced various approaches. One such approach are the s-step Krylov subspace methods (also called "communication-avoiding" Krylov subspace methods); for a thorough treatment of background, related work, and performance experiments, see, e.g., the theses [5, 27]. In s-step Krylov subspace methods, instead of performing one iteration at a time, the iterations are performed in blocks of s; i.e., in each iteration, the Krylov subspace is expanded by $O(s)$ dimensions by computing $O(s)$ new basis vectors and then all inner products between the new basis vectors needed for the next s iterations are computed in one block operation. In this way, computing the inner products for s iterations only requires a single global synchronization, decreasing the synchronization cost per iteration by a factor of $O(s)$. This approach has been shown to lead to significant speedups for a number of problems and real-world applications (see, e.g., [35, 43]). In the remainder of the paper, we will refer to the matrices whose columns consist of the $O(s)$ basis vectors computed in each block as *s-step basis matrices*. Further details of the s-step CG method are discussed in section 2. We emphasize that our use of the overloaded term "s-step methods" here differs from other works, e.g., [14, 30] and [21, §9.2.7], in which 's-step method' refers to a type of restarted Lanczos procedure.

In exact arithmetic the s-step CG method produces the exact same iterates as the classical CG method, but their behavior can differ significantly in finite precision. In both s-step and classical Krylov subspace methods, rounding errors due to finite precision arithmetic have two basic effects: a decrease in attainable accuracy and a delay of convergence. It has long been known that for s-step Krylov subspace methods, as s is increased (and so the condition numbers of the s-step basis matrices increase), the attainable accuracy decreases and the convergence delay increases relative to the classical CG method (see, e.g., [10]). At the extreme, if the parameter s is chosen to be too large, the $O(s)$-dimensional bases computed for each block can be numerically rank deficient and the s-step method can fail to converge.

This sensitive numerical behavior poses a practical obstacle to optimizing the performance of s-step methods, and diminishes their usability and reliability. In a setting where the performance of CG is communication-bound, we expect that up to some point, increasing s will decrease the time per iteration. If we pick s only based on minimizing the time per iteration, however, we can run into problems. First, the finite precision error may cause a large convergence delay, negating any potential performance gain with respect to the overall runtime. Since the number of iterations required for convergence for a given s value is not known a priori, choosing the s value that results in the fastest time-to-solution is a difficult problem. Second, the chosen s parameter may cause s-step CG to fail to converge to the user-specified accuracy. In this case, the particular problem is *unsolvable* by the s-step CG method.

Requiring the user to choose the parameter s thus diminishes the practical usability of s-step Krylov subspace methods. It is therefore imperative that we develop a better understanding of the convergence rate and accuracy in finite precision s-step CG and other s-step Krylov subspace methods. Our hope is that by studying the theoretical properties of methods designed for large-scale computations in finite precision, we can develop methods and techniques that are efficient, capable of meeting application-dependent accuracy constraints, and which do not require that the user have extensive knowledge of numerical linear algebra.

Toward this goal, in this paper we develop a variable s-step CG method in which the s values can vary between blocks and are chosen automatically by the algorithm

such that a user-specified accuracy is attained. We call this the *adaptive s-step CG method*. The condition for determining s in each block is derived using a bound on the size of the gap between the true residuals $b - Ax_i$ and the updated residuals r_i computed in finite precision. We show that in order to prevent a loss of attainable accuracy, in a given block of iterations, the condition number of the s-step basis matrix should be no larger than a quantity related to the inverse of the 2-norms of the residuals computed within that block. In other words, as the method converges, more ill-conditioned bases (i.e., larger s values) can be used without a loss of attainable accuracy. Furthermore, we show that the adaptive s-step CG method can be implemented in a way that does not increase the synchronization cost per block versus (fixed) s-step CG.

Whereas the fixed s-step CG method may not attain the same accuracy as classical CG, our numerical experiments on a set of example matrices demonstrate that our adaptive s-step CG method can achieve the same accuracy as classical CG, while still reducing the number of synchronizations required by up to over a factor of 4.

In section 2 we give a brief background on s-step Krylov subspace methods and detail the s-step CG method, and in section 3 we discuss related work. Section 4 presents our main contribution; in section 4.1, we derive a theoretical bound on how large the s-step basis matrix condition number can be in terms of the norms of the updated residuals without affecting the attainable accuracy, and in section 4.2, we detail the adaptive s-step CG algorithm that makes use of this bound in determining the number of iterations to perform in each block. Numerical experiments for a small set of test matrices are presented in section 5. We discuss possible extensions and future work and briefly conclude in section 6.

2. Background: s-step Krylov subspace methods. The basic idea behind s-step Krylov subspace methods is to grow the underlying Krylov subspace $O(s)$ dimensions at a time and perform the necessary orthogonalization with a single global synchronization. Within a block of s iterations, the vector updates are performed implicitly by updating the coordinates of the iteration vectors in the $O(s)$-dimensional basis spanned by the columns of the s-step basis matrix. The algorithmic details vary depending on the particular Krylov subspace method considered, but the general approach is the same. There is a wealth of literature related to s-step Krylov subspace methods. We give a brief summary of early related work below; a thorough bibliography can be found in [27, Table 1.1].

The s-step approach was first used in the context of Krylov subspace methods by Van Rosendale [41], who developed a variant of the parallel conjugate gradient method which minimizes inner product data dependencies with the goal of exposing more parallelism. The term "s-step Krylov subspace methods" was first used by Chronopoulos and Gear, who developed an s-step CG method [9, 10]. This work was followed in subsequent years by the development of other s-step variants of Orthomin and GMRES [11], Arnoldi and Symmetric Lanczos [31, 32], MINRES, GCR, and Orthomin [8, 13], Nonsymmetric Lanczos [33], and Orthodir [12]. Walker similarly used s-step bases, but with the goal of improving stability in GMRES by replacing the modified Gram-Schmidt orthogonalization process with Householder QR [42].

Many of the earliest s-step Krylov subspace methods used monomial bases (i.e., $[v, Av, A^2v, \ldots]$) for the computed s-step basis matrices, but it was found that convergence often could not be guaranteed for $s > 5$ (see, e.g., [10]). This motivated research into the use of other more well-conditioned bases for the Krylov subspaces, including scaled monomial bases [26], Chebyshev bases [29, 16, 17], and Newton bases [1, 20].

The growing cost of communication in large-scale sparse problems has created a recent resurgence of interest in the implementation, optimization, and development of s-step Krylov subspace methods; see, e.g., the recent works [18, 27, 35, 46, 23, 34, 44, 45, 43].

2.1. The s-step conjugate gradient method. In this work, we focus on the s-step conjugate gradient method (Algorithm 2), which is equivalent to the s-step CG method (Algorithm 2) in exact arithmetic. In the remainder of this section, we give a brief explanation of the s-step CG method; further details on the derivation, implementation, and performance of the s-step CG method can be found in, e.g., [5, 27].

Algorithm 1 Conjugate gradient (CG)

Input: $n \times n$ symmetric positive definite matrix A, length-n vector b, and initial approximation x_1 to $Ax = b$

Output: Approximate solution x_{i+1} to $Ax = b$ with updated residual r_{i+1}

1: $r_1 = b - Ax_1$, $p_1 = r_1$
2: **for** $i = 1, 2, \ldots,$ until convergence **do**
3: $\quad \alpha_i = r_i^T r_i / p_i^T A p_i$
4: $\quad q_i = \alpha_i p_i$
5: $\quad x_{i+1} = x_i + q_i$
6: $\quad r_{i+1} = r_i - A q_i$
7: $\quad \beta_i = r_{i+1}^T r_{i+1} / r_i^T r_i$
8: $\quad p_{i+1} = r_{i+1} + \beta_i p_i$
9: **end for**

The s-step CG method consists of an outer loop, indexed by k, which iterates over the blocks of iterations, and an inner loop, which iterates over iterations $j \in \{1, \ldots, s\}$ within a block. For clarity, we globally index iterations by $i \equiv sk + j$. It follows from the properties of CG that for $j \in \{0, \ldots, s\}$,

$$\begin{aligned} &p_{sk+j+1}, r_{sk+j+1} \in \mathcal{K}_{s+1}(A, p_{sk+1}) + \mathcal{K}_s(A, r_{sk+1}) \quad \text{and} \\ &x_{sk+j+1} - x_{sk+1} \in \mathcal{K}_s(A, p_{sk+1}) + \mathcal{K}_{s-1}(A, r_{sk+1}). \end{aligned} \tag{2.1}$$

Then the CG vectors for the next s iterations lie in the sum of the column spaces of the matrices

$$\begin{aligned} &\mathcal{P}_{k,s} = [\rho_0(A)p_{sk+1}, \ldots, \rho_s(A)p_{sk+1}], \text{ with } \operatorname{span}(\mathcal{P}_{k,s}) = \mathcal{K}_{s+1}(A, p_{sk+1}) \quad \text{and} \\ &\mathcal{R}_{k,s} = [\rho_0(A)r_{sk+1}, \ldots, \rho_{s-1}(A)r_{sk+1}], \text{ with } \operatorname{span}(\mathcal{R}_{k,s}) = \mathcal{K}_s(A, r_{sk+1}), \end{aligned} \tag{2.2}$$

where $\rho_j(z)$ is a polynomial of degree j satisfying the three-term recurrence

$$\begin{aligned} &\rho_0(z) = 1, \qquad \rho_1(z) = (z - \theta_0)\rho_0(z)/\gamma_0, \quad \text{and} \\ &\rho_j(z) = \left((z - \theta_{j-1})\rho_{j-1}(z) - \sigma_{j-2}\rho_{j-2}(z)\right)/\gamma_{j-1}. \end{aligned} \tag{2.3}$$

We define the s-step basis matrix $\mathcal{Y}_{k,s} = [\mathcal{P}_{k,s}, \mathcal{R}_{k,s}]$ and define $\underline{\mathcal{Y}}_{k,s}$ to be the same as $\mathcal{Y}_{k,s}$ except with all zeros in columns $s+1$ and $2s+1$. Note that the s's in the subscripts are unnecessary here, but will be useful in describing the variable s-step CG method later on. Under certain assumptions on the nonzero structure of A, $\mathcal{Y}_{k,s}$ can be computed in parallel in each outer loop for the same asymptotic latency cost

as a single matrix-vector product using the so-called "matrix powers kernel" (see [18] for details). Since the columns of $\mathcal{Y}_{k,s}$ satisfy (2.3), we can write

$$A\underline{\mathcal{Y}}_{k,s} = \mathcal{Y}_{k,s}\mathcal{B}_{k,s}, \tag{2.4}$$

where $\mathcal{B}_{k,s}$ is a $(2s+1) \times (2s+1)$ block tridiagonal matrix defined by the 3-term recurrence coefficients in (2.3).

By (2.1), there exist vectors $p'_{k,j+1}$, $r'_{k,j+1}$, and $x'_{k,j+1}$ that represent the coordinates of the CG iterates p_{sk+j+1}, r_{sk+j+1}, and $x_{sk+j+1} - x_{sk+1}$, respectively, in $\mathcal{Y}_{k,s}$ for $j \in \{0, \ldots, s\}$. That is,

$$\begin{aligned} p_{sk+j+1} &= \mathcal{Y}_{k,s}p'_{k,j+1}, \\ r_{sk+j+1} &= \mathcal{Y}_{k,s}r'_{k,j+1}, \quad \text{and} \\ x_{sk+j+1} - x_{sk+1} &= \mathcal{Y}_{k,s}x'_{k,j+1}. \end{aligned} \tag{2.5}$$

Therefore, in the inner loop of s-step CG, rather than update the CG vectors explicitly, we instead update their coordinates in $\mathcal{Y}_{k,s}$, i.e., for $j \in \{1, \ldots, s\}$,

$$\begin{aligned} x'_{k,j+1} = x'_{k,j} + \alpha_{sk+j}p'_{k,j}, \quad r'_{k,j+1} = r'_{k,j} - \alpha_{sk+j}\mathcal{B}_{k,s}p'_{k,j}, \\ \text{and} \quad p'_{k,j+1} = r'_{k,j+1} + \beta_{sk+j}p'_{k,j}. \end{aligned}$$

Thus the matrix-vector product with A in each iteration becomes a matrix-vector product with the much smaller matrix $\mathcal{B}_{k,s}$. This along with the length-$(2s+1)$ vector updates can be performed locally by each processor in each inner loop iteration.

We can also reduce communication in computing the inner products. Letting $G_{k,s} = \mathcal{Y}_{k,s}^T\mathcal{Y}_{k,s}$, and using (2.5) and (2.4), α_{sk+j} and β_{sk+j} can be computed by

$$\begin{aligned} \alpha_{sk+j} &= \left(r'^T_{k,j}G_{k,s}r'_{k,j}\right) / \left(p'^T_{k,j}G_{k,s}\mathcal{B}_{k,s}p'_{k,j}\right) \quad \text{and} \\ \beta_{sk+j} &= \left(r'^T_{k,j+1}G_{k,s}r'_{k,j+1}\right) / \left(r'^T_{k,j}G_{k,s}r'_{k,j}\right). \end{aligned}$$

The matrix $G_{k,s}$ can be computed with one global synchronization per outer loop iteration, which, in terms of asymptotic parallel latency, costs the same as a single inner product. As $\mathcal{B}_{k,s}$ and $G_{k,s}$ are dimension $(2s+1) \times (2s+1)$, α_{sk+j} and β_{sk+j} can be computed locally by each processor within the inner loop.

3. Related work. In this section, we discussed related work in the areas of variable s-step Krylov subspace methods, the analysis of maximum attainable accuracy in finite precision CG, and inexact Krylov subspace methods.

3.1. Variable s-step Krylov subspace methods. Williams et al. [43] use a variable s-step BICGSTAB as the coarse grid solver in a multigrid method. The technique used here, termed "telescoping", is motivated by the observation than some coarse grid solves converge after only a few iterations, whereas other solves take significantly longer. By starting with a small s value and gradually increasing it, they ensure that s-step BICGSTAB with larger s is only used when the solve takes enough iterations to amortize the additional costs associated with s-step CG.

Imberti and Erhel [28] have recently developed an s-step GMRES method that allows variable s sizes. They recommend choosing s_k values according to the Fibonacci sequence up to some maximum value $s_{\max}$, i.e., starting with $s_0 = s_1 = 1$ and for $k \geq 2$, $s_k = \min(s_{\max}, s_{k-1} + s_{k-2})$. In this way, the sequence $\{s_i\}$ used by Imberti and Erhel is predetermined.

Algorithm 2 s-step conjugate gradient

Input: $n \times n$ symmetric positive definite matrix A, length-n vector b, and initial approximation x_1 to $Ax = b$

Output: Approximate solution x_{sk+s+1} to $Ax = b$ with updated residual r_{sk+s+1}

$r_1 = b - Ax_1$, $p_1 = r_1$
for $k = 0, 1, \ldots,$ until convergence **do**
  Compute $\mathcal{Y}_{k,s} = [\mathcal{P}_{k,s}, \mathcal{R}_{k,s}]$ according to (2.2)
  Compute $G_{k,s} = \mathcal{Y}_{k,s}^T \mathcal{Y}_{k,s}$
  Assemble $\mathcal{B}_{k,s}$ such that (2.4) holds
  $p'_{k,1} = [1, 0_{1,2s}]^T$, $r'_{k,1} = [0_{1,s+1}, 1, 0_{1,s-1}]^T$, $x'_{k,1} = [0_{1,2s+1}]^T$
  for $j = 1$ to s **do**
    $\alpha_{sk+j} = \left(r'^T_{k,j} G_{k,s} r'_{k,j}\right) / \left(p'^T_{k,j} G_{k,s} \mathcal{B}_{k,s} p'_{k,j}\right)$
    $q'_{k,j} = \alpha_{sk+j} p'_{k,j}$
    $x'_{k,j+1} = x'_{k,j} + q'_{k,j}$
    $r'_{k,j+1} = r'_{k,j} - \mathcal{B}_{k,s} q'_{k,j})$
    $\beta_{sk+j} = \left(r'^T_{k,j+1} G_{k,s} r'_{k,j+1}\right) / \left(r'^T_{k,j} G_{k,s} r'_{k,j}\right)$
    $p'_{k,j+1} = r'_{k,j+1} + \beta_{sk+j} p'_{k,j}$
  end for
  Recover iterates $\{p_{sk+s+1}, r_{sk+s+1}, x_{sk+s+1}\}$ according to (2.5)
end for

In contrast, in our approach, the sequence $\{s_i\}$ is dynamically chosen based on the 2-norm of the updated residual (which is not necessarily monotonically decreasing in CG). In addition, our method is designed such that a user-specified accuracy ε^* can be attained when $\varepsilon^* \geq \varepsilon_{\text{CG}}$, where ε_{CG} is the accuracy attained by classical CG for the same problem.

3.2. Maximum attainable accuracy in Krylov subspace methods. In both s-step and classical variants of Krylov subspace methods, finite precision roundoff error in updates to the approximate solution x_i and the residual r_i in each iteration can cause the updated residual r_i and the true residual $b - Ax_i$ to grow further and further apart as the iterations proceed. If this deviation grows large, it can limit the *maximum attainable accuracy*, i.e., the accuracy with which we can solve $Ax = b$ on a computer with unit round-off ε. Analyses of maximum attainable accuracy in CG and other classical KSMs are given by Greenbaum [22], van der Vorst and Ye [40], Sleijpen, van der Vorst, and Fokkema [37], Sleijpen, van der Vorst, and Modersitzki [38], Björck, Elfving, and Strakoš [2], and Gutknecht and Strakoš [24]. One important result of these analyses is the insight that loss of accuracy can be caused at a very early stage of the computation, which can not be corrected in later iterations. Analyses of the maximum attainable accuracy in s-step CG and the s-step biconjugate gradient method (BICG) can be found in [6, 5].

3.3. Inexact Krylov subspace methods. The term "inexact Krylov subspace methods" refers to Krylov subspace methods in which the matrix-vector products Av are computed inexactly as $(A + E)v$ where E represents some error matrix (due to either finite precision computation or intentional approximation). It is assumed that all other computations (orthogonalization and vector updates) are performed exactly. Using an analysis of the deviation of the true and updated residuals, it is shown that under these assumptions the size of the perturbation term E can be allowed to

grow inversely proportional to the norm of the updated residual without adversely affecting the attainable accuracy. This theory has potential application in mixed-precision Krylov subspace methods, as well as in situations where the operator A is expensive to apply and can be approximated in order to gain performance.

Much work has been dedicated to inexact Krylov subspace methods. In [36], Simoncini and Szyld provide a general theory for inexact variants applicable to both symmetric and nonsymmetric linear systems and eigenvalue problems and use this to derive criteria for bounding the size of E in such a way that does not diminish the attainable accuracy of the method. A similar analysis was given by van de Eschof and Sleijpen [39]. These analyses confirm and improve upon the earlier empirical results for the inexact GMRES method of Bouras and Fraysseé [3] and the inexact conjugate gradient method of Bouras, Frayssé, and Giraud [4].

Our work is very closely related to the theory of inexact Krylov subspace methods in that our analysis results in a somewhat analogous "relaxation strategy", where instead of relaxing the accuracy of the matrix-vector products we are instead relaxing a constraint on the condition numbers of the computed s-step basis matrices. In addition, our analysis assumes that all parts of the computation are performed in a fixed precision ε.

Algorithm 3 Variable s-step conjugate gradient

Input: $n \times n$ symmetric positive definite matrix A, length-n vector b, initial approximation x_1 to $Ax = b$, sequence of s values $(s_0, s_1, \ldots)$

Output: Approximate solution x_{m+s_k+1} to $Ax = b$ with updated residual r_{m+s_k+1}

$r_1 = b - Ax_1$, $p_1 = r_1$
$m = 0$
for $k = 0, 1, \ldots,$ until convergence **do**
  Compute $\mathcal{Y}_{k,s_k} = [\mathcal{P}_{k,s_k}, \mathcal{R}_{k,s_k}]$ according to (2.2)
  Compute $G_{k,s_k} = \mathcal{Y}_{k,s_k}^T \mathcal{Y}_{k,s_k}$
  Assemble $\mathcal{B}_{k,s_k}$ such that (2.4) holds
  $p'_{k,1} = [1,\, 0_{1,2s_k}]^T$, $r'_{k,1} = [0_{1,s_k+1},\, 1,\, 0_{1,s_k-1}]^T$, $x'_{k,1} = [0_{1,2s_k+1}]^T$
  for $j = 1$ to s_k **do**
    $\alpha_{m+j} = \left(r'^T_{k,j} G_{k,s_k} r'_{k,j}\right) / \left(p'^T_{k,j} G_{k,s_k} \mathcal{B}_{k,s_k} p'_{k,j}\right)$
    $q'_{k,j} = \alpha_{m+j} p'_{k,j}$
    $x'_{k,j+1} = x'_{k,j} + q'_{k,j}$
    $r'_{k,j+1} = r'_{k,j} - \mathcal{B}_{k,s_k} q'_{k,j})$
    $\beta_{m+j} = \left(r'^T_{k,j+1} G_{k,s_k} r'_{k,j+1}\right) / \left(r'^T_{k,j} G_{k,s_k} r'_{k,j}\right)$
    $p'_{k,j+1} = r'_{k,j+1} + \beta_{m+j} p'_{k,j}$
  end for
  Recover iterates $\{p_{m+s_k+1}, r_{m+s_k+1}, x_{m+s_k+1}\}$ according to (2.5)
  $m = m + s_k$
end for

4. Variable s-step CG. As discussed, the conditioning of the s-step basis matrices in s-step Krylov subspace methods affects the rate of convergence and attainable accuracy. In the s-step CG method (Algorithm 2), the Krylov subspace basis is computed in blocks of size s in each outer loop. This method can be generalized to allow the blocks to be of varying size; in other words, a different s value can be used in each outer loop iteration. We denote the block size in outer iteration k by s_k, where $1 \leq s_k \leq s_{\max}$ for some maximum s value $s_{\max}$. We call this a *variable s-step Krylov*

subspace method. A general variable s-step CG method is shown in Algorithm 3.

In this section, we derive a variable s-step CG method which automatically determines s_k in each outer loop such that a user-specified accuracy can be attained, which we call the adaptive s-step CG method. Our analysis is based on the maximum attainable accuracy analysis for s-step CG in [6], which shows that the attainable accuracy of s-step CG depends on the condition numbers of the s-step basis matrices $\mathcal{Y}_{k,s}$ computed in each outer loop. In summary, we show that if in outer loop k, s_k is selected such that the condition number of the basis matrix $\mathcal{Y}_{k,s_k}$ is inversely proportional to the maximum 2-norm of the residual vectors computed within outer loop k, then the approximate solution produced by the variable s-step CG method can be as accurate as the approximate solution produced by the classical CG method. Further, our method allows for the user to specify the desired accuracy, so in the case that a less accurate solution is required, our method automatically adjusts to allow for higher s_k values. This effectively exploits the tradeoff between accuracy and performance in s-step Krylov subspace methods.

In this section we derive our adaptive s-step approach, which stems from a bound on the growth in outer loop k of the gap between the true and updated residuals in finite precision.

4.1. A constraint on basis matrix condition number. In the remainder of the paper, hats denote quantities computed in finite precision, ε denotes the unit round-off, $\|\cdot\|$ denotes the 2-norm, and $\kappa(A)$ denotes the 2-norm condition number $\|A^{-1}\|\|A\|$. To simplify the analysis, we drop all $O(\varepsilon^2)$ terms (and higher order terms in ε). Writing the true residual $b - A\widehat{x}_i = b - A\widehat{x}_i - \widehat{r}_i + \widehat{r}_i$, we can write the bound

$$\|b - A\widehat{x}_i\| \leq \|b - A\widehat{x}_i - \widehat{r}_i\| + \|\widehat{r}_i\|.$$

Then as $\|\widehat{r}_i\| \to 0$ (i.e., as the method converges), we have $\|b - A\widehat{x}_i\| \approx \|b - A\widehat{x}_i - \widehat{r}_i\|$. The size of the *residual gap*, i.e., $\|b - A\widehat{x}_i - \widehat{r}_i\|$, therefore determines the attainable accuracy of the method (see the related references in section 3.2).

We begin by considering the rounding errors made in the variable s-step CG method (Algorithm 3). In outer loop k and inner loops $j \in \{1, \ldots, s_k\}$ of finite precision variable s-step CG, using standard models of floating point error (e.g., [21, §2.4]) we have

$$A\underline{\widehat{\mathcal{Y}}}_{k,s_k} = \widehat{\mathcal{Y}}_{k,s_k}\mathcal{B}_{k,s_k} - E_k, \tag{4.1}$$
$$\text{where} \quad \|E_k\| \leq \varepsilon\left((3+N_A)\|A\|\|\underline{\widehat{\mathcal{Y}}}_{k,s_k}\| + 4\|\widehat{\mathcal{Y}}_{k,s_k}\|\|\mathcal{B}_{k,s_k}\|\right),$$
$$\widehat{x}'_{k,j+1} = \widehat{x}'_{k,j} + \widehat{q}'_{k,j} + \xi_{k,j+1}, \tag{4.2}$$
$$\text{where} \quad \|\xi_{k,j+1}\| \leq \varepsilon\|\widehat{x}'_{k,j+1}\|,$$
$$\widehat{r}'_{k,j+1} = \widehat{r}'_{k,j} - \mathcal{B}_{k,s_k}\widehat{q}'_{k,j} + \eta_{k,j+1}, \tag{4.3}$$
$$\text{where} \quad \|\eta_{k,j+1}\| \leq \varepsilon(\|\widehat{r}'_{k,j+1}\| + 3\|\mathcal{B}_{k,s_k}\|\|\widehat{q}'_{k,j}\|),$$
$$\widehat{x}_{m+j+1} = \widehat{\mathcal{Y}}_{k,s_k}\widehat{x}'_{k,j+1} + \widehat{x}_{m+1} + \phi_{m+j+1}, \tag{4.4}$$
$$\text{where} \quad \|\phi_{m+j+1}\| \leq \varepsilon(\|\widehat{x}_{m+j+1}\| + 3j\|\widehat{\mathcal{Y}}_{k,s_k}\|\|\widehat{x}'_{k,j+1}\|), \text{ and}$$
$$\widehat{r}_{m+j+1} = \widehat{\mathcal{Y}}_{k,s_k}\widehat{r}'_{k,j+1} + \psi_{m+j+1},, \tag{4.5}$$
$$\text{where} \quad \|\psi_{m+j+1}\| \leq 3j\varepsilon\|\widehat{\mathcal{Y}}_{k,s_k}\|\|\widehat{r}'_{k,j+1}\|,$$

where N_A denotes the maximum number of nonzeros per row in A and $m = \sum_{i=0}^{k-1} s_i$. Let $\delta_i \equiv b - A\widehat{x}_i - \widehat{r}_i$ denote the residual gap in iteration i. Similar to the analysis

in [6], for variable s-step CG we can write the growth of the residual gap in iteration $m+j+1$ in outer loop k as

$$\delta_{m+j+1} - \delta_{m+1} = -\sum_{i=1}^{j} \left(A\underline{\hat{\mathcal{Y}}}_{k,s_k} \xi_{k,i+1} + \widehat{\mathcal{Y}}_{k,s_k} \eta_{k,i+1} - E_k \widehat{q}'_{k,i} \right) - A\phi_{m+j+1} - \psi_{m+j+1}.$$

Our goal is now to manipulate a bound on this quantity in order to derive a method for determining the largest s_k value we can use in outer loop k such that the desired accuracy is attained. Using standard techniques, we have the norm-wise bound

$$\begin{aligned}
\|\delta_{m+j+1} - \delta_{m+1}\| \leq{}& \varepsilon\|A\|\|\underline{\hat{\mathcal{Y}}}_{k,s_k}\| \sum_{i=1}^{j} \|\widehat{x}'_{k,i+1}\| + \varepsilon\|\widehat{\mathcal{Y}}_{k,s_k}\| \sum_{i=1}^{j} \|\widehat{r}'_{k,i+1}\| \\
&+ \varepsilon \left((3+N_A)\|A\|\|\underline{\hat{\mathcal{Y}}}_{k,s_k}\| + 7\|\widehat{\mathcal{Y}}_{k,s_k}\|\|\mathcal{B}_{k,s_k}\| \right) \sum_{i=1}^{j} \|\widehat{q}'_{k,i}\| \\
&+ \varepsilon\|A\|\|\widehat{x}_{m+j+1}\| + 3j\varepsilon\|A\|\|\widehat{\mathcal{Y}}_{k,s_k}\|\|\widehat{x}'_{k,j+1}\| + 3j\varepsilon\|\widehat{\mathcal{Y}}_{k,s_k}\|\|\widehat{r}'_{k,j+1}\| \\
\leq{}& \varepsilon \left((7+2N_A)\|A\|\|\underline{\hat{\mathcal{Y}}}_{k,s_k}\| + 14\|\widehat{\mathcal{Y}}_{k,s_k}\|\|\mathcal{B}_{k,s_k}\| \right) \sum_{i=1}^{j} \|\widehat{x}'_{k,i+1}\| \\
&+ \varepsilon\|\widehat{\mathcal{Y}}_{k,s_k}\| \sum_{i=1}^{j} \|\widehat{r}'_{k,i+1}\| \\
&+ \varepsilon\|A\|\|\widehat{x}_{m+j+1}\| + 3j\varepsilon\|A\|\|\widehat{\mathcal{Y}}_{k,s_k}\|\|\widehat{x}'_{k,j+1}\| + 3j\varepsilon\|\widehat{\mathcal{Y}}_{k,s_k}\|\|\widehat{r}'_{k,j+1}\|,
\end{aligned}$$

where we have used $\sum_{i=1}^{j} \|\widehat{q}'_{k,i}\| \leq (2+\varepsilon) \sum_{i=1}^{j} \|\widehat{x}'_{k,i+1}\|$. To simplify the notation, we let $\tau_k = \|\mathcal{B}_{k,s_k}\|/\|A\|$ and $N_k = 7 + 2N_A + 14\tau_k$. We note that $\|\mathcal{B}_{k,s_k}\|$ depends on the polynomial basis used and is not expected to be too large; e.g., for the monomial basis, $\|\mathcal{B}_{k,s_k}\| = 1$. The above bound can then be written

$$\begin{aligned}
\|\delta_{m+j+1} - \delta_{m+1}\| \leq{}& \varepsilon N_k\|A\|\|\widehat{\mathcal{Y}}_{k,s_k}\| \sum_{i=1}^{j} \|\widehat{x}'_{k,i+1}\| + \varepsilon\|\widehat{\mathcal{Y}}_{k,s_k}\| \sum_{i=1}^{j} \|\widehat{r}'_{k,i+1}\| \\
&+ \varepsilon\|A\|\|\widehat{x}_{m+j+1}\| + 3j\varepsilon\|A\|\|\widehat{\mathcal{Y}}_{k,s_k}\|\|\widehat{x}'_{k,j+1}\| + 3j\varepsilon\|\widehat{\mathcal{Y}}_{k,s_k}\|\|\widehat{r}'_{k,j+1}\|.
\end{aligned} \tag{4.6}$$

Note that assuming $\widehat{\mathcal{Y}}_{k,s_k}$ is full rank, we have

$$\begin{aligned}
\|\widehat{\mathcal{Y}}_{k,s_k}\|\|\widehat{r}'_{k,i+1}\| &= \|\widehat{\mathcal{Y}}_{k,s_k}\|\|\widehat{\mathcal{Y}}^{+}_{k,s_k} \widehat{\mathcal{Y}}_{k,s_k} \widehat{r}'_{k,i+1}\| \\
&\leq \kappa(\widehat{\mathcal{Y}}_{k,s_k})\|\widehat{\mathcal{Y}}_{k,s_k} \widehat{r}_{m+i+1} - \psi_{m+i+1}\| \\
&= \kappa(\widehat{\mathcal{Y}}_{k,s_k})\|\widehat{r}_{m+i+1}\| + O(\varepsilon), \quad \text{and,} \qquad (4.7) \\
\|\widehat{\mathcal{Y}}_{k,s_k}\|\|\widehat{x}'_{k,i+1}\| &= \|\widehat{\mathcal{Y}}_{k,s_k}\|\|\widehat{\mathcal{Y}}^{+}_{k,s_k} \widehat{\mathcal{Y}}_{k,s_k} \widehat{x}'_{k,i+1}\| \\
&\leq \kappa(\widehat{\mathcal{Y}}_{k,s_k})\|\widehat{\mathcal{Y}}_{k,s_k} \widehat{x}'_{k,i+1}\|. \qquad (4.8) \\
& \qquad (4.9)
\end{aligned}$$

To bound $\|\widehat{\mathcal{Y}}_{k,s_k} \widehat{x}'_{k,i+1}\|$ in terms of the size of the residuals, notice that by (4.4),

$$\widehat{\mathcal{Y}}_{k,s_k} \widehat{x}'_{k,i+1} = \widehat{x}_{m+i+1} - \widehat{x}_{m+1} - \phi_{m+i+1}$$

$$
\begin{aligned}
&= (x - \widehat{x}_{m+1}) - (x - \widehat{x}_{m+i+1}) - \phi_{m+i+1} \\
&= A^{-1}(b - A\widehat{x}_{m+1}) - A^{-1}(b - A\widehat{x}_{m+i+1}) - \phi_{m+i+1} \\
&= A^{-1}(\widehat{r}_{m+1} + \delta_{m+1}) - A^{-1}(\widehat{r}_{m+i+1} + \delta_{m+i+1}) - \phi_{m+i+1} \\
&= A^{-1}(\widehat{r}_{m+1} - \widehat{r}_{m+i+1}) + A^{-1}(\delta_{m+1} - \delta_{m+i+1}) - \phi_{m+i+1},
\end{aligned}
$$

so

$$
\|\widehat{\mathcal{Y}}_{k,s_k}\widehat{x}'_{k,i+1}\| \leq \|A^{-1}\|(\|\widehat{r}_{m+1}\| + \|\widehat{r}_{m+i+1}\|) + O(\varepsilon).
$$

Thus together with (4.8) we can write

$$
\|\widehat{\mathcal{Y}}_{k,s_k}\|\|\widehat{x}'_{k,i+1}\| \leq \kappa(\widehat{\mathcal{Y}}_{k,s_k})\|A^{-1}\|(\|\widehat{r}_{m+1}\| + \|\widehat{r}_{m+i+1}\|) + O(\varepsilon).
$$

This allows us to write the bound (4.6) in the form

$$
\begin{aligned}
\|\delta_{m+j+1} - \delta_{m+1}\| \leq\ & 2\varepsilon N_k \kappa(A)\kappa(\widehat{\mathcal{Y}}_{k,s_k}) \sum_{i=1}^{j} \|\widehat{r}_{m+i+1}\| \\
&+ \varepsilon \left(N_k j \kappa(A)\kappa(\widehat{\mathcal{Y}}_{k,s_k}) + 3j\kappa(A)\kappa(\widehat{\mathcal{Y}}_{k,s_k}) \right) \|\widehat{r}_{m+1}\| \\
&+ \varepsilon \left(3j\kappa(A)\kappa(\widehat{\mathcal{Y}}_{k,s_k}) + 3j\kappa(A)\kappa(\widehat{\mathcal{Y}}_{k,s_k}) \right) \|\widehat{r}_{m+j+1}\| \\
&+ \varepsilon\|A\|\|\widehat{x}_{m+j+1}\|.
\end{aligned}
$$

Using $x - \widehat{x}_{m+j+1} = A^{-1}\widehat{r}_{m+j+1} + \delta_{m+j+1}$, we have

$$
\|\widehat{x}_{m+j+1} - x\| \leq \|A^{-1}\|\|\widehat{r}_{m+j+1}\| + O(\varepsilon),
$$

and writing $\widehat{x}_{m+j+1} = \widehat{x}_{m+j+1} - x + x$, we have

$$
\begin{aligned}
\|\delta_{m+j+1} - \delta_{m+1}\| \leq\ & 2N_k\varepsilon\kappa(A)\kappa(\widehat{\mathcal{Y}}_{k,s_k}) \left(j(\|\widehat{r}_{m+1}\| + \|\widehat{r}_{m+j+1}\|) + \sum_{i=1}^{j} \|\widehat{r}_{m+i+1}\| \right) \\
&+ \varepsilon\|A\|\|x\|.
\end{aligned}
$$

Then

$$
\|\delta_{m+j+1} - \delta_{m+1}\| \leq 6N_k j\varepsilon\kappa(A)\kappa(\widehat{\mathcal{Y}}_{k,s_k}) \left(\max_{\ell\in\{0,\ldots,j\}} \|\widehat{r}_{m+\ell+1}\| \right) + \varepsilon\|A\|\|x\|,
$$

and letting $c_{k,j} = 6N_k j\kappa(A)$, this bound can be written

$$
\|\delta_{m+j+1} - \delta_{m+1}\| \leq c_{k,j}\varepsilon\kappa(\widehat{\mathcal{Y}}_{k,s_k}) \left(\max_{\ell\in\{0,\ldots,j\}} \|\widehat{r}_{m+\ell+1}\| \right) + \varepsilon\|A\|\|x\|. \tag{4.10}
$$

This gives a norm-wise bound on the growth of the residual gap due to finite precision errors made in outer iteration k.

Letting $m_\ell = \sum_{i=0}^{\ell-1} s_i$, notice that we have

$$
\|\delta_{m_k+s_k+1}\| \leq \|\delta_1\| + \sum_{\ell=0}^{k} \|\delta_{m_\ell+s_\ell+1} - \delta_{m_\ell+1}\|.
$$

Suppose that at the end of the iterations we want the size of the residual gap to be on the order ε^* (where $\varepsilon^* \gtrsim O(\varepsilon)\|A\|\|x\|$, as this is the best we could expect). Then assuming the number of outer iterations k is not too high, this can be accomplished by requiring that in each outer loop, for $j \in \{1, \ldots, s_k\}$,

$$\|\delta_{m+j+1} - \delta_{m+1}\| \leq \varepsilon^*.$$

From (4.10), this means that we must have, for all inner iterations $j \in \{1, \ldots, s_k\}$,

$$\kappa(\widehat{\mathcal{Y}}_{k,s_k}) \leq \frac{\varepsilon^*}{c_{k,j}\varepsilon \max_{\ell \in \{0,\ldots,j\}} \|\widehat{r}_{m+\ell+1}\|}. \tag{4.11}$$

The bound (4.11) tells us that as the 2-norm of the residual decreases, we can tolerate a more ill-conditioned basis matrix $\widehat{\mathcal{Y}}_{k,s}$ without detriment to the attainable accuracy. Since $\kappa(\widehat{\mathcal{Y}}_{k,s})$ grows with increasing s_k, this suggests that s_k can be allowed to increase as the method converges at a rate proportional to the inverse of the residual 2-norm. This naturally gives a relaxation strategy that is somewhat analogous to those derived for inexact Krylov subspace methods (see section 3.3). We discuss the implementation of this strategy in the following subsection.

4.2. The adaptive s-step CG method. In this section, we propose a variable s-step CG method that makes use of the constraint (4.11) in determining s_k in each outer loop k, which we call the *adaptive* s-step CG method. Our approach uses the largest s_k value possible in each outer iteration k such that (4.11) is satisfied up to some $s_{\max}$ set by the user (e.g., $s_{\max}$ could be selected by auto-tuning to find the s value that minimizes the time per iteration in s-step CG). Since a variable s-step CG method produces the same iterates as classical CG when $s_k = 1$ for all outer loops k, we expect that our adaptive s-step CG method can attain accuracy ε^* when $\varepsilon^* \geq \varepsilon_{\text{CG}}$, where ε_{CG} is the attainable accuracy of classical CG. When the right-hand-side of (4.11) is less than one (due to, e.g., ε^* being set too small), the inequality (4.11) cannot be satisfied. In this case, we default to using $s_k = 1$ (i.e., we perform an iteration of classical CG).

In CG, it is the A-norm of the error rather than the 2-norm of the residual that is minimized, which means that in some inner loop iteration j, we can have $\|\widehat{r}_{m+j+1}\| > \|\widehat{r}_{m+1}\|$. Since $\|\widehat{r}_{m+i+1}\|$ for $i > 1$ are not known at the start of the outer loop iteration, we must determine s_k online, i.e., during the iterations. We therefore first construct a basis matrix that satisfies (4.11) under the assumption that $\|\widehat{r}_{m+i+1}\| \leq \|\widehat{r}_{m+1}\|$. Within each inner loop, we can perform an inexpensive check to ensure that (4.11) is still satisfied after computing the next residual; if not, we break from the inner loop, discard the current basis matrix, and begin a new outer loop iteration. We elaborate on this approach in the remainder of this section. An optimized high-performance implementation of our variable s-step approach is outside the scope of the present work, although we will briefly discuss how one might implement this strategy efficiently such that each outer loop still requires only a single global synchronization and the number of wasted matrix-vector products is kept to a minimum.

The bound (4.11) indicates that when $\widehat{\mathcal{Y}}_{k,s_k}$ is constructed at the beginning of outer loop k, we should enforce

$$\kappa(\widehat{\mathcal{Y}}_{k,s_k}) \leq \frac{\varepsilon^*}{c_{k,s_{\max}}\varepsilon\|\widehat{r}_{m+1}\|}. \tag{4.12}$$

Therefore we first construct an $\bar{s}_k$-step basis matrix $\widehat{\mathcal{Y}}_{k,\bar{s}_k}$ using some $\bar{s}_k \leq s_{\max}$, compute the Gram matrix $\widehat{G}_{k,\bar{s}_k} = \widehat{\mathcal{Y}}_{k,\bar{s}_k}^T \widehat{\mathcal{Y}}_{k,\bar{s}_k}$, and then use the condition numbers of the leading principle submatrices of the blocks of $\widehat{G}_{k,\bar{s}_k}$ to compute the condition numbers of the leading columns of the blocks of $\widehat{\mathcal{Y}}_{k,\bar{s}_k}$. To be concrete, for $i \in \{1, 2, \ldots, \bar{s}_k\}$, we have (using MATLAB notation)

$$\widehat{\mathcal{Y}}_{k,i} = \left[\widehat{\mathcal{Y}}_{k,\bar{s}_k}(:, 1{:}i{+}1), \widehat{\mathcal{Y}}_{k,\bar{s}_k}(:, i_s{:}i_s{+}i{-}1)\right]$$

and

$$\widehat{G}_{k,i} = \begin{bmatrix} \widehat{G}_{k,\bar{s}_k}(1{:}i{+}1, 1{:}i{+}1) & \widehat{G}_{k,\bar{s}_k}(1{:}i{+}1, i_s{:}i_s{+}i{-}1) \\ \widehat{G}_{k,\bar{s}_k}(i_s{:}i_s{+}i{-}1, 1{:}i{+}1) & \widehat{G}_{k,\bar{s}_k}(i_s{:}i_s{+}i{-}1, i_s{:}i_s{+}i{-}1) \end{bmatrix},$$

where $i_s = \bar{s}_k + 2$. Then since $\sqrt{\kappa(\widehat{G}_{k,i})} = \kappa(\widehat{\mathcal{Y}}_{k,i})$, we can easily compute the condition number of the i-step basis matrix $\widehat{\mathcal{Y}}_{k,i}$, i.e., the basis matrix needed to perform i inner loop iterations, for $i \in \{1, \ldots, \bar{s}_k\}$. We therefore let $\tilde{s}_k \in \{1, \ldots, \bar{s}_k\}$ be the largest value possible such that

$$\sqrt{\kappa(\widehat{G}_{k,\tilde{s}_k})} \leq \frac{\varepsilon^*}{c_{k,s_{\max}} \varepsilon \|\widehat{r}_{m+1}\|} \tag{4.13}$$

holds.

To handle the case where $\|\widehat{r}_{m+j+1}\| > \|\widehat{r}_{m+1}\|$ for some inner iteration j, we can add a check within the inner loop iterations: after computing $\widehat{r}'_{k,j+1}$, we check if

$$\kappa(\widehat{\mathcal{Y}}_{k,\tilde{s}_k}) \leq \frac{\varepsilon^*}{c_{k,j} \varepsilon \|\widehat{r}_{m+j+1}\|} \approx \frac{\varepsilon^*}{c_{k,j} \varepsilon \left(\widehat{r}'^T_{k,j+1} \widehat{G}_{k,\tilde{s}_k} \widehat{r}'_{k,j+1}\right)^{1/2}} \tag{4.14}$$

is still satisfied. Note that since $\|\widehat{r}_{m+j+1}\| = \left(\widehat{r}'^T_{k,j+1} \widehat{G}_{k,\tilde{s}_k} \widehat{r}'_{k,j+1}\right)^{1/2} + O(\varepsilon)$, performing this check does not require any additional communication. If (4.14) is satisfied for all $j \in \{1, \ldots, \tilde{s}_k\}$, we will perform all $s_k = \tilde{s}_k$ inner loop iterations. If this is not satisfied for some j, we break from the current inner loop iteration and begin the next outer loop; i.e., we only perform $s_k = j$ inner loop iterations in outer loop k.

Although performing the checks for determining when to break from the current outer loop do not incur any additional communication cost in terms of the number of messages sent between processors, there is, however, potential wasted effort in terms of computation and the number of words moved when $\bar{s}_k > \tilde{s}_k$ and/or $\tilde{s}_k > s_k$, since in this case we have computed more basis vectors than we need to perform s_k inner loop updates. One way to mitigate this is to only allow $\bar{s}_k$ to grow by at most f vectors in each outer loop, i.e., $\bar{s}_k \leq s_{k-1} + f \leq s_{\max}$, where f is some small positive integer (e.g., $f = 1$ or $f = 2$).

The approach outlined above is summarized in Algorithm 4. Note that in Algorithm 4, we have specified that the function $c_{k,j}$ can be input by the user. Because our bounds are not tight, in many cases using $c_{k,j} = 6N_k j\kappa(A)$ (from the bound (4.11)) overestimates the error, which results in the use of smaller s_k values than necessary.

Note that although there is no expectation that $\|\widehat{r}_{m+j+1}\|$ decreases between iterations in CG, as long as the method is converging, we still expect the residual norm to trend downward. Therefore one could make the simplifying assumption that

$\|\widehat{r}_{m+1}\| \geq \|\widehat{r}_{m+j+1}\|$ for $j \in \{1, \ldots, \tilde{s}_k\}$. In this case, the value $s_k = \tilde{s}_k$ can be determined at the beginning of outer loop k before performing any inner iterations, and so the checks (4.14) are unnecessary. This simplification comes at the potential cost of reduced reliability of the adaptive s-step method, although in our experimenting we have observed that taking $s_k = \tilde{s}_k$ usually works well even when the residual norm is nonmonotonic.

Because our strategy is based on a loose upper bound on the deviation of residuals, we do not require particularly accurate estimates of $\kappa(\widehat{\mathcal{Y}}_{k,i})$. Therefore alternative methods for inexpensively approximating this quantity could also be used in practice.

Algorithm 4 Adaptive s-step conjugate gradient

Input: $n \times n$ symmetric positive definite matrix A, length-n vector b, initial approximation x_1 to $Ax = b$, maximum s value $s_{\max}$, initial s value $\bar{s}_0$, maximum basis growth factor f, desired convergence tolerance ε^*, function $c_{k,j}$

Output: Approximate solution x_{m+s_k+1} to $Ax = b$ with updated residual r_{m+s_k+1}

$r_1 = b - Ax_1$, $p_1 = r_1$
$m = 0$
for $k = 0, 1, \ldots,$ until convergence **do**
  if $k \neq 0$ **then**
    $\bar{s}_k = \min(s_{k-1} + f, s_{\max})$
  end if
  Compute $\bar{s}_k$-step basis matrix $\mathcal{Y}_{k,\bar{s}_k} = [\mathcal{P}_{k,\bar{s}_k}, \mathcal{R}_{k,\bar{s}_k}]$ according to (2.2)
  Compute $G_{k,\bar{s}_k} = \mathcal{Y}_{k,\bar{s}_k}^T \mathcal{Y}_{k,\bar{s}_k}$
  Determine $\tilde{s}_k$ by (4.13); assemble $\mathcal{Y}_{k,\tilde{s}_k}$ and $G_{k,\tilde{s}_k}$
  Store estimate $\gamma \approx \kappa(\widehat{\mathcal{Y}}_{k,\tilde{s}_k})$
  Assemble $\mathcal{B}_{k,\tilde{s}_k}$ such that (2.4) holds
  $p'_{k,1} = [1,\, 0_{1,2\tilde{s}_k}]^T$, $r'_{k,1} = [0_{1,\tilde{s}_k+1},\, 1,\, 0_{1,\tilde{s}_k-1}]^T$, $x'_{k,1} = [0_{1,2\tilde{s}_k+1}]^T$
  for $j = 1$ to $\tilde{s}_k$ **do**
    $s_k = j$
    $\alpha_{m+j} = \left(r'^T_{k,j} G_{k,\tilde{s}_k} r'_{k,j}\right) / \left(p'^T_{k,j} G_{k,\tilde{s}_k} \mathcal{B}_{k,\tilde{s}_k} p'_{k,j}\right)$
    $q'_{k,j} = \alpha_{m+j} p'_{k,j}$
    $x'_{k,j+1} = x'_{k,j} + q'_{k,j}$
    $r'_{k,j+1} = r'_{k,j} - \mathcal{B}_{k,\tilde{s}_k} q'_{k,j})$
    $\beta_{m+j} = \left(r'^T_{k,j+1} G_{k,\tilde{s}_k} r'_{k,j+1}\right) / \left(r'^T_{k,j} G_{k,\tilde{s}_k} r'_{k,j}\right)$
    $p'_{k,j+1} = r'_{k,j+1} + \beta_{m+j} p'_{k,j}$
    if $\gamma \geq \varepsilon^* / c_{k,j} \varepsilon \left(\widehat{r}'^T_{k,j+1} G_{k,\tilde{s}_k} \widehat{r}'_{k,j+1}\right)^{1/2}$ **then**
      break from inner loop
    end if
  end for
  Recover iterates $\{p_{m+s_k+1}, r_{m+s_k+1}, x_{m+s_k+1}\}$ according to (2.5)
  $m = m + s_k$
end for

5. Numerical experiments. In this section, we test the adaptive s-step CG method (Algorithm 4) on a number of small examples from the University of Florida Sparse Matrix Collection [15] using MATLAB. The test examples were chosen to exemplify various convergence trajectories of CG, which allows us to demonstrate the validity of our relaxation condition in which s_k depends on the rate of convergence.

The test problems used here are small, so we would not use (fixed or variable) s-step methods here in practice; nevertheless, they still serve to exhibit the numerical behavior of the methods.

All tests were run in double precision, i.e., $\varepsilon \approx 2^{-53}$, and use a right-hand side b with entries $1/\sqrt{n}$ where n is the dimension of A. We use matrix equilibration, i.e., $A \leftarrow D^{-1/2}AD^{-1/2}$ where D is a diagonal matrix of the largest entries in each row of A. Properties of the test matrices (post-equilibration) are shown in Table 5.1.

Our results are plotted in Figures 5.1-5.5. For each test matrix, we ran classical CG (Algorithm 1), as well as both (fixed) s-step CG (Algorithm 2) and adaptive s-step CG (Algorithm 4) for $s_{\max}$ values 4, 8, and 10 (for fixed s-step CG, $s = s_{\max}$). For each matrix and $s_{\max}$ value, we tested two different ε^* values: $\varepsilon^* = \varepsilon_{\text{CG}}$, the maximum attainable accuracy of classical CG, and $\varepsilon^* = 1e-6$. The monomial basis was used for both variable and fixed s-step methods in all tests (note that there is no technical restriction to the monomial basis; the analysis used in deriving the adaptive s-step CG method still applies regardless of the polynomials used in constructing $\widehat{\mathcal{Y}}_{k,\bar{s}_k}$). In all tests for adaptive s-step CG, we set $\bar{s}_0 = f = s_{\max}$ (see Algorithm 4), as this gives the optimal choice of s_k in each iteration. (We also tried using smaller values of f, e.g., $f = 1$ and $f = 2$; this did not significantly change the total number of outer loop iterations performed). We used $c_{k,j} = 1$ in our condition for determining the number of inner loop iterations to perform in all experiments except for tests with matrices bcsstk09 (Figure 5.4) and ex5 (Figure 5.5); for these problems, convergence is more irregular, and thus the bound (4.11) is tighter.

Each figure has a corresponding table (Tables 5.2-5.6) which lists the number of outer loop iterations (a proxy for the latency cost or number of global synchronizations) required for convergence of the true residual 2-norm to level ε^*. For classical CG, the number reported gives the total number of iterations. Dashes in the tables indicate that the true residual diverged or stagnated before reaching level ε^*.

TABLE 5.1
Test Matrix Properties.

Matrix	n	nnz	$\|A\|$	$\kappa(A)$
gr_30_30	900	7744	1.5	1.9e2
mesh3e1	289	1377	1.8	8.6e0
nos6	675	3255	2.0	3.5e6
bcsstk09	1083	18437	2.0	1.0e4
ex5	27	279	3.8	5.7e7

5.1. Observations. In all experiments, the adaptive s-step CG method was able to converge to the desired tolerance ε^*, whereas for $\varepsilon^* = \varepsilon_{\text{CG}}$ and higher s values ($s = 8$ and/or $s = 10$), fixed s-step CG may not achieve convergence to this level. We point out that even when the matrix is very well-conditioned, choosing s too large can very negatively affect the convergence rate in fixed s-step CG; see, e.g., the case with $s = 10$ and $\varepsilon^* = 1e-6$ in Figure 5.2, where s-step CG requires more than 8 times the number of synchronizations than classical CG. In tests where both fixed s-step CG and adaptive s-step CG converge, adaptive s-step CG takes at most 2 additional outer iterations versus fixed s-step CG (although in most cases it takes as many or fewer outer iterations). We note that in all cases, variable s-step CG uses fewer outer loop iterations (synchronization points) than classical CG, so we still expect a performance advantage on latency-bound problems.

The experiments using $\varepsilon^* = 1e{-}6$ demonstrate the ability of the adaptive s-step

CG method to automatically adjust to use the largest s_k values possible (according to our bound) that will eventually attain the desired accuracy. For example, in Figure 5.1 for gr_30_30, for all tests using $\varepsilon^* = 1e{-}6$, adaptive s-step CG uses $s_k = s_{\max}$ in every outer iteration, defaulting to the fixed s-step method. In contrast, in plots on the left using $\varepsilon^* = \varepsilon_{\text{CG}}$, adaptive s-step CG uses smaller s_k values at the beginning and then increases up to $s_k = s_{\max}$ and is able to attain accuracy ε_{CG} whereas fixed s-step CG fails when for $s = 8$ and $s = 10$.

Below we list the sequence of s_k values chosen in adaptive s-step CG for tests with $\varepsilon^* = \varepsilon_{\text{CG}}$ and $s_{\max} = 10$ (plots in the lower left in each figure):

- gr_30_30: 1, 1, 2, 2, 2, 3, 3, 3, 4, 5, 6, 8, 10, 10
- mesh3e1: 1, 1, 2, 4, 6, 9, 10
- nos6: 6, 1, 2, 3, 4, 5, 4, 4, 1, 4, 4, 5, 6, 7, 8, 10, 10, 10, 10, 10, 10, 10, 10, 10, 10, 10, 10, 10, 10, 10, 10, 10, 10, 10
- bcsstk09: 2, 2, 3, 3, 3, 3, 3, 1, 2, 3, 3, 3, 2, 2, 2, 2, 3, 3, 3, 3, 4, 4, 4, 4, 5, 5, 6, 7, 1, 6, 6, 7, 7, 7, 7, 8, 8, 8, 9, 9, 10, 10, 10, 10, 10, 10, 10, 10, 10, 10
- ex5: 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 3, 1, 1, 2, 1, 1, 1, 3, 3, 2, 1, 1, 1, 2, 1, 1, 2, 4, 6, 2, 4, 5, 7, 2, 6, 1, 7, 6

There are a few things to notice. First, as expected, the sequence of s_k values is monotonically increasing when the residual 2-norm decreases monotonically or close to it (matrices gr_30_30 and mesh3e1), but when the residual 2-norm oscillates, we no longer have monotonicity of the s_k values (matrices nos6, bcsstk09, and ex5). Second, it may be that $s_{\max}$ is never reached, as is the case for matrix ex5. It is also worth pointing out that the sequence of s_k values need not start with $s_0 = 1$; for nos6, $s_0 = 6$ and for bcsstk09, $s_0 = 2$.

It may be that on a particular machine for a particular matrix structure, it is not beneficial in terms of minimizing time per iteration to go above a small value of $s_{\max}$, say $s_{\max} = 4$. In this case, as long as the matrix is reasonably well-conditioned and the accuracy required is not too high, the fixed s-step CG method can probably be used without too much trouble. However, even in this case the adaptive s-step CG method can provide reliability while still maintaining the reduction in synchronizations of the fixed s-step method, at the cost of potentially wasting a few matrix-vector multiplications in each outer iteration.

Looking at the number of outer loop iterations in Tables 5.2-5.6, we see that in our tests for adaptive s-step CG it is often the case that increasing the $s_{\max}$ value past a certain point gives diminishing returns in terms of reducing synchronization. This may represent a fundamental limit of the possible performance using the monomial basis for these problems. In these cases it may be that better performance (higher s_k values) can be achieved only by using a better-conditioned polynomial basis for constructing $\widehat{\mathcal{Y}}_{k,\bar{s}_k}$ (e.g., Newton or Chebyshev bases).

6. Conclusions and future work. In this work, we developed the adaptive s-step CG method, which is a variable s-step CG method where the parameter s_k is determined automatically in each outer loop such that a user-specified accuracy is attainable. The method for determining s_k is based on a bound on the residual gap within outer loop k, from which we derive a constraint on the condition number of the s-step basis matrix generated in outer loop k. The computations required for determining s_k can be performed without introducing any additional synchronizations in each outer loop. Our numerical experiments demonstrate that the adaptive s-step CG method is able to attain up to the same accuracy as classical CG while still significantly reducing the number of outer loop iterations (i.e., synchronizations)

performed. In contrast, the fixed s-step CG method can fail to converge to this accuracy when $s > 4$.

Additionally, the adaptive s-step CG method adjusts to use larger s_k when an application requires a less accurate solution, automatically optimizing the tradeoff between speed and accuracy. In this way, a user can use the matrix structure and machine parameters to estimate the value $s_{\max}$ that will minimize the time per iteration (e.g., by guessing, or offline auto-tuning), input the accuracy they require, and the adaptive s-step CG method will choose s_k as large as it can (i.e., reducing as much synchronization as it can) while still converging to the desired level. This takes the burden of parameter selection off the user. Our adaptive s-step CG method thus makes significantly advances in improving reliability and usability in s-step Krylov subspace methods.

Our future work includes a high-performance implementation and performance experiments in order to compare classical CG, fixed s-step CG, and adaptive s-step CG on large-scale problems. Another potential improvement is to obtain a tighter bound to replace (4.10), which could avoid the requirement of adjusting the function $c_{k,j}$ (which represents a multiplicative factor in the bounds (4.10) and (4.11)) in Algorithm 4.

The adaptive s-step CG method is designed to maximize s_k subject to a constraint on accuracy, but obviously it would also be useful to optimize s_k to achieve the *overall* lowest latency/synchronization cost. This would require a better understanding of the finite precision convergence behavior of both s-step and classical Krylov subspace methods, which is an active area of research.

We believe that this adaptive approach based on a bound on the residual gap can be extended to other s-step Krylov subspace methods as well. The same approach outlined here can be used in s-step BICG (see, e.g., [7]), which uses the same recurrences for the solution and updated residual. We conjecture that a similar approach applies to s-step variants of other recursively computed residual methods (see [22]), like s-step BICGSTAB [7], as well as methods like s-step GMRES. In the case of s-step GMRES, it may be possible to simplify the adaptive algorithm: since the residual norm monotonically decreases, the sequence of s_k values should monotonically increase. Such a sequence of s_k values was used in the variable s-step GMRES method of Imberti and Erhel [28].

As a final comment, we note that in our analysis in section 4, we used equation (4.10) to determine a bound on $\kappa(\widehat{\mathcal{Y}}_{k,s_k})$ that guarantees accuracy on the level ε^*. We could have just as easily used equation (4.10) to develop a mixed precision s-step CG method where s is fixed, but the precision used in computing in outer loop k varies. In other words, rearranging the inequality (4.10) we see that, given that we've computed an s-step basis with condition number $\kappa(\widehat{\mathcal{Y}}_{k,s})$, if we want to achieve accuracy ε^*, we must have that

$$\varepsilon_k \leq \frac{\varepsilon^*}{c_{k,j}\kappa(\widehat{\mathcal{Y}}_{k,s}) \max_{\ell\in\{0,\ldots,j\}} \|\widehat{r}_{m+\ell+1}\|}.$$

Investigating the potential for mixed precision s-step Krylov subspace methods based on this analysis remains future work.

FIG. 5.1. *Convergence of the 2-norm of the true residual for the matrix gr_30_30. Plots on the left use ε^* =3.4e-14 and plots on the right use ε^* =1e-6. Plots are shown for three s (or s_{max}) values: 4 (top), 8 (middle), and 10 (bottom). Iterations where communication takes place are noted with markers (red circles for s-step CG, blue stars for variable s-step CG, and black dots for classical CG). The horizontal dashed line shows the requested accuracy ε^*.*

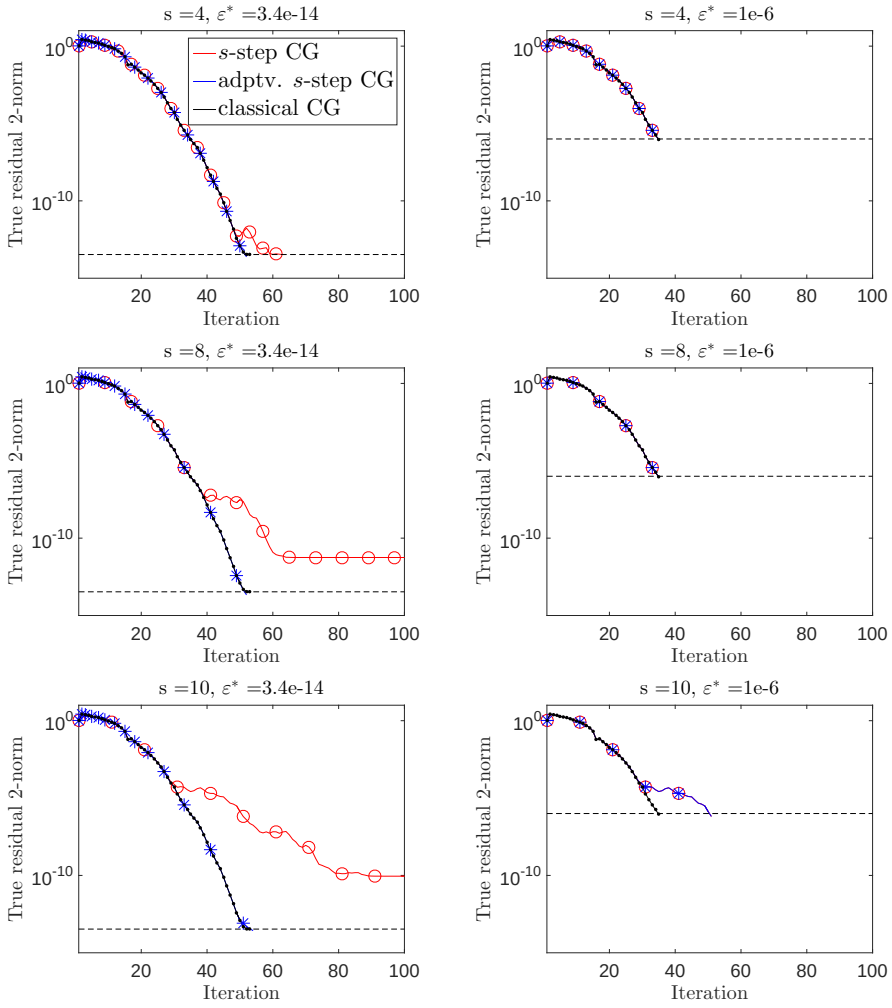

TABLE 5.2

Number of global synchronizations (outer loop iterations) in tests for matrix gr_30_30 (corresponding to Figure 5.1). Dashes (-) in the table indicate that the method failed to converge to the desired level ε^*.

		fixed s-step CG	adaptive s-step CG	classical CG
ε^* =3.4e-14	$s = 4$	16	17	52
	$s = 8$	-	14	
	$s = 10$	-	14	
ε^* =1e-6	$s = 4$	9	9	34
	$s = 8$	5	5	
	$s = 10$	5	5	

FIG. 5.2. *Convergence of the 2-norm of the true residual for the matrix mesh3e1. Plots on the left use ε^* =1.0e-14 and plots on the right use ε^* =1e-6. Plots are shown for three s (or s_{max}) values: 4 (top), 8 (middle), and 10 (bottom). Iterations where communication takes place are noted with markers (red circles for s-step CG, blue stars for variable s-step CG, and black dots for classical CG). The horizontal dashed line shows the requested accuracy ε^*.*

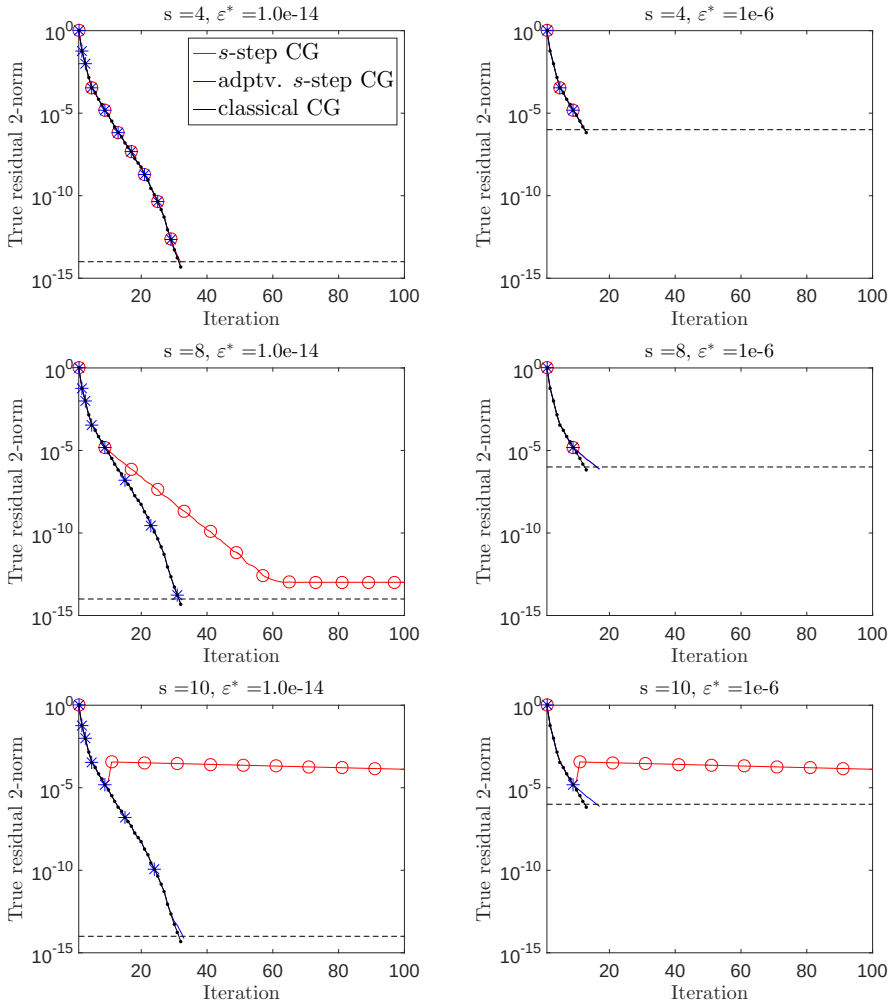

TABLE 5.3
Number of global synchronizations (outer loop iterations) in tests for matrix mesh3e1 (corresponding to Figure 5.2). Dashes (-) in the table indicate that the method failed to converge to the desired level ε^*.

		fixed s-step CG	adaptive s-step CG	classical CG
ε^* =1e-14	$s = 4$	8	10	31
	$s = 8$	-	8	
	$s = 10$	-	7	
ε^* =1e-6	$s = 4$	3	3	12
	$s = 8$	2	2	
	$s = 10$	99	2	

FIG. 5.3. *Convergence of the 2-norm of the true residual for the matrix nos6. Plots on the left use ε^* =5.5e-10 and plots on the right use ε^* =1e-6. Plots are shown for three s (or s_{max}) values: 4 (top), 8 (middle), and 10 (bottom). Iterations where communication takes place are noted with markers (red circles for s-step CG, blue stars for variable s-step CG, and black dots for classical CG). The horizontal dashed line shows the requested accuracy ε^*.*

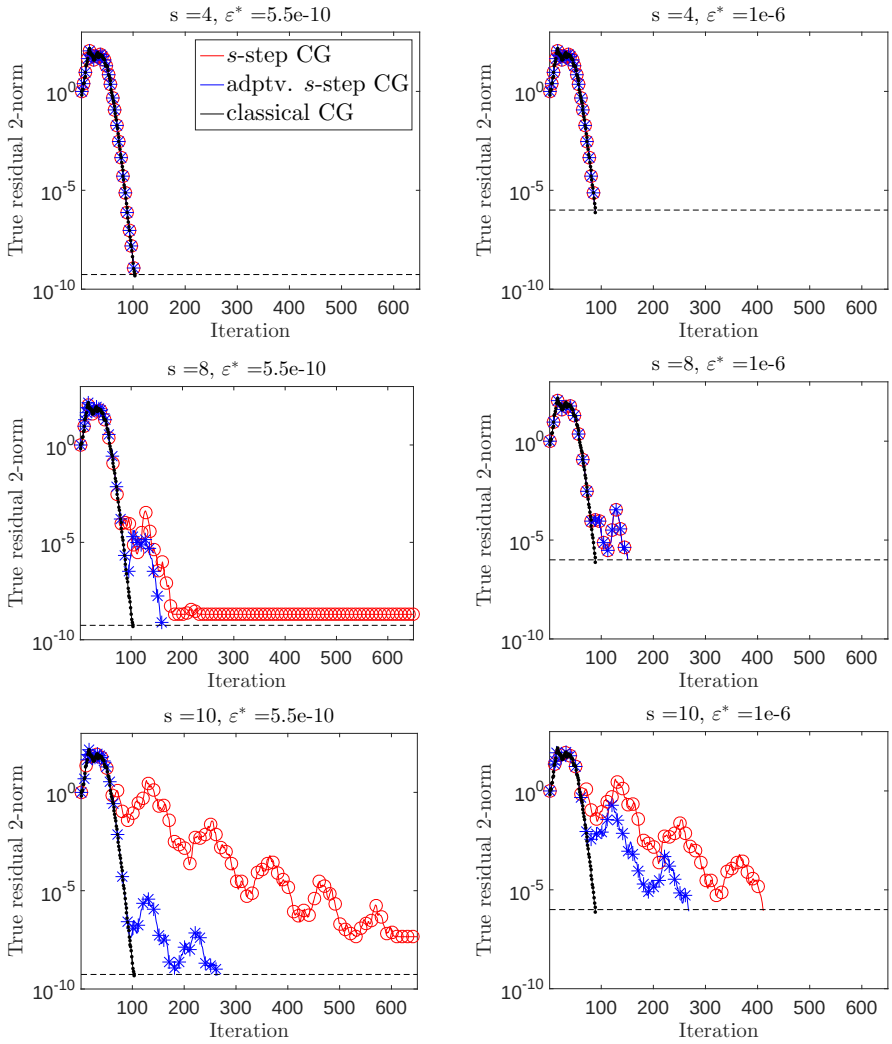

TABLE 5.4

Number of global synchronizations (outer loop iterations) in tests for matrix nos6 (corresponding to Figure 5.3). Dashes (-) in the table indicate that the method failed to converge to the desired level ε^*.

		fixed s-step CG	adaptive s-step CG	classical CG
ε^* =5.5e-10	$s = 4$	26	26	103
	$s = 8$	-	29	
	$s = 10$	-	36	
ε^* =1e-6	$s = 4$	22	22	88
	$s = 8$	19	19	
	$s = 10$	41	29	

FIG. 5.4. *Convergence of the 2-norm of the true residual for the matrix bcsstk09. Plots on the left use ε^* =1.6e-12 and plots on the right use ε^* =1e-6. Plots are shown for three s (or s_{max}) values:* 4 *(top),* 8 *(middle), and* 10 *(bottom). Iterations where communication takes place are noted with markers (red circles for s-step CG, blue stars for variable s-step CG, and black dots for classical CG). The horizontal dashed line shows the requested accuracy ε^*.*

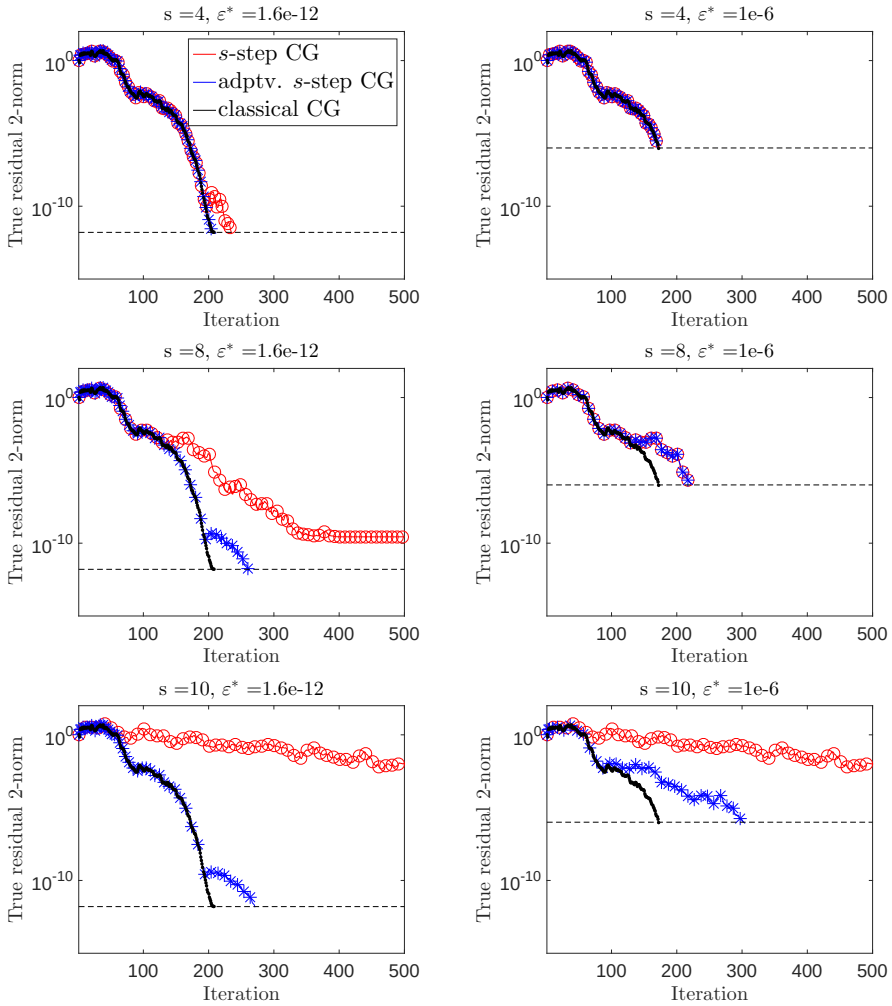

TABLE 5.5

Number of global synchronizations (outer loop iterations) in tests for matrix bcsstk09 (corresponding to Figure 5.4). Dashes (-) in the table indicate that the method failed to converge to the desired level ε^*. Tests used $c_{k,j} = 10$.

		fixed s-step CG	adaptive s-step CG	classical CG
ε^* =1.6e-12	$s = 4$	59	59	207
	$s = 8$	-	52	
	$s = 10$	-	50	
ε^* =1e-6	$s = 4$	43	43	171
	$s = 8$	28	28	
	$s = 10$	122	33	

FIG. 5.5. *Convergence of the 2-norm of the true residual for the matrix ex5. Plots on the left use ε^* =1e-8 and plots on the right use ε^* =1e-6. Plots are shown for three s (or s_{max}) values: 4 (top), 8 (middle), and 10 (bottom). Iterations where communication takes place are noted with markers (red circles for s-step CG, blue stars for variable s-step CG, and black dots for classical CG). The horizontal dashed line shows the requested accuracy ε^*.*

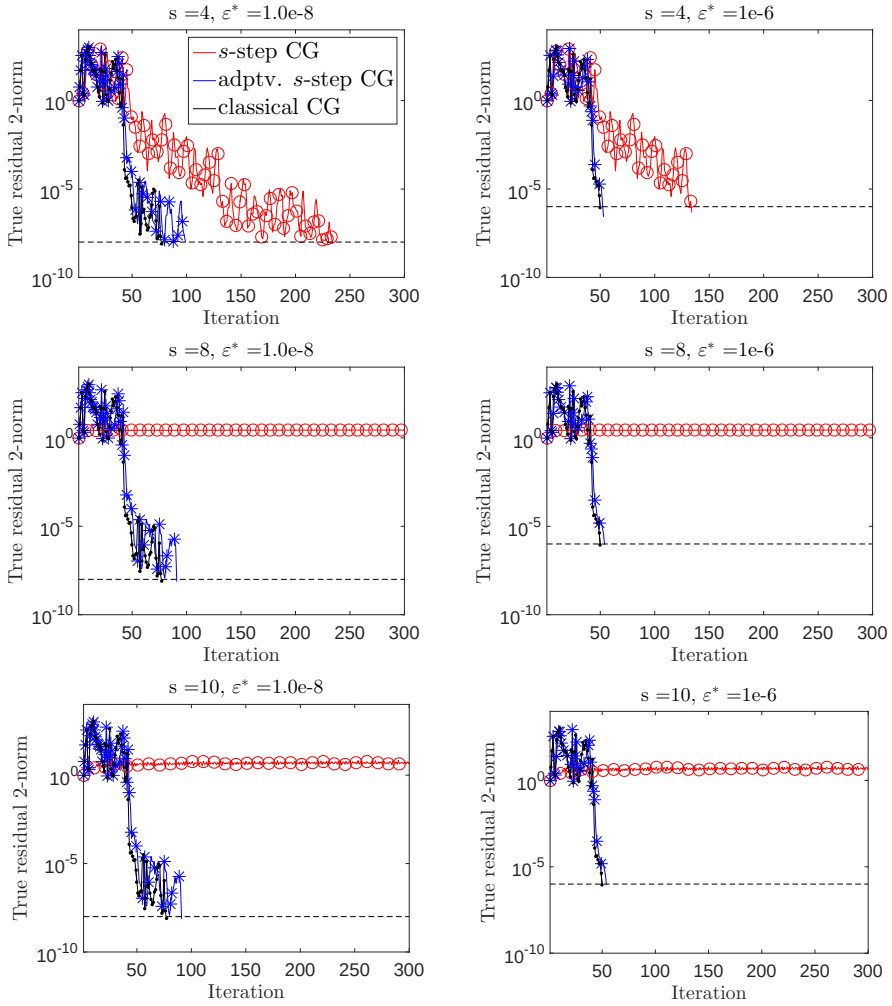

TABLE 5.6

Number of global synchronizations (outer loop iterations) in tests for matrix ex5 (corresponding to Figure 5.5). Dashes (-) in the table indicate that the method failed to converge to the desired level ε^*. Tests used $c_{k,j} = \sqrt{\kappa(A)}$.

		fixed s-step CG	adaptive s-step CG	classical CG
ε^* =1e-8	$s = 4$	59	48	76
	$s = 8$	-	45	
	$s = 10$	-	45	
ε^* =1e-6	$s = 4$	34	27	49
	$s = 8$	-	27	
	$s = 10$	-	27	

REFERENCES

[1] Z. Bai, D. Hu, and L. Reichel, *A Newton basis GMRES implementation*, IMA J. Numer. Anal., 14 (1994), pp. 563–581.

[2] Å. Björck, T. Elfving, and Z. Strakoš, *Stability of conjugate gradient and Lanczos methods for linear least squares problems*, SIAM J. Matrix Anal. Appl., 19 (1998), pp. 720–736.

[3] A. Bouras and V. Frayssé, *A relaxation strategy for inexact matrix-vector products for Krylov methods*, Tech. Report CERFACS TR/PA/00/15, European Centre for Research and Advanced Training in Scientific Computation, 2000.

[4] A. Bouras, V. Frayssé, and L. Giraud, *A relaxation strategy for inner-outer linear solvers in domain decomposition methods*, Tech. Report CERFACS TR/PA/00/17, European Centre for Research and Advanced Training in Scientific Computation, 2000.

[5] E. Carson, *Communication-avoiding Krylov subspace methods in theory and practice*, PhD thesis, EECS Dept., U.C. Berkeley, Aug 2015.

[6] E. Carson and J. Demmel, *A residual replacement strategy for improving the maximum attainable accuracy of s-step Krylov subspace methods*, SIAM J. Matrix Anal. Appl., 35 (2014), pp. 22–43.

[7] E. Carson, N. Knight, and J. Demmel, *Avoiding communication in nonsymmetric Lanczos-based Krylov subspace methods*, SIAM J. Sci. Comp., 35 (2013).

[8] A. Chronopoulos, *s-step iterative methods for (non)symmetric (in)definite linear systems*, SIAM J. Numer. Anal., 28 (1991), pp. 1776–1789.

[9] A. Chronopoulos and C. Gear, *On the efficient implementation of preconditioned s-step conjugate gradient methods on multiprocessors with memory hierarchy*, Parallel Comput., 11 (1989), pp. 37–53.

[10] A. Chronopoulos and C. Gear, *s-step iterative methods for symmetric linear systems*, J. Comput. Appl. Math, 25 (1989), pp. 153–168.

[11] A. Chronopoulos and S. Kim, *s-step Orthomin and GMRES implemented on parallel computers*, Tech. Report 90/43R, USMI, Minneapolis, MN, 1990.

[12] A. Chronopoulos and D. Kincaid, *On the Odir iterative method for non-symmetric indefinite linear systems*, Numer. Lin. Alg. Appl., 8 (2001), pp. 71–82.

[13] A. Chronopoulos and C. Swanson, *Parallel iterative s-step methods for unsymmetric linear systems*, Parallel Comput., 22 (1996), pp. 623–641.

[14] J. Cullum and W. Donath, *A block Lanczos algorithm for computing the q algebraically largest eigenvalues and a corresponding eigenspace of large, sparse, real symmetric matrices*, in Proc. of the 1974 IEEE Conf. on Decision and Control, IEEE, 1974, pp. 505–509.

[15] T. Davis and Y. Hu, *The University of Florida sparse matrix collection*, ACM Trans. Math. Software, 38 (2011), pp. 1–25.

[16] E. de Sturler, *A parallel variant of GMRES(m)*, in Proceedings of the 13th IMACS World Congress on Computational and Applied Mathematics. IMACS, Criterion Press, vol. 9, 1991.

[17] E. De Sturler and H. van der Vorst, *Reducing the effect of global communication in GMRES(m) and CG on parallel distributed memory computers*, Appl. Numer. Math., 18 (1995), pp. 441–459.

[18] J. Demmel, M. Hoemmen, M. Mohiyuddin, and K. Yelick, *Avoiding communication in computing Krylov subspaces*, Tech. Report UCB/EECS-2007-123, EECS Dept., U.C. Berkeley, Oct 2007.

[19] J. Dongarra et al., *The international exascale software project roadmap*, Int. J. High Perform. Comput. Appl., 25 (2011), pp. 3–60.

[20] J. Erhel, *A parallel GMRES version for general sparse matrices*, Electronic Transactions on Numerical Analysis, 3 (1995), pp. 160–176.

[21] G. Golub and C. Van Loan, *Matrix computations*, Johns Hopkins University Press, Baltimore, MD, 3 ed., 1996.

[22] A. Greenbaum, *Estimating the attainable accuracy of recursively computed residual methods*, SIAM J. Matrix Anal. Appl., 18 (1997), pp. 535–551.

[23] L. Grigori and S. Moufawad, *Communication avoiding ILU(0) preconditioner*, Rapport de recherche RR-8266, INRIA, Mar. 2013, `http://hal.inria.fr/hal-00803250`.

[24] M. Gutknecht and Z. Strakoš, *Accuracy of two three-term and three two-term recurrences for Krylov space solvers*, SIAM J. Matrix Anal. Appl., 22 (2000), pp. 213–229.

[25] M. Hestenes and E. Stiefel, *Methods of conjugate gradients for solving linear systems*, J. Nat. Bur. Standards, 49 (1952), pp. 409–436.

[26] A. Hindmarsh and H. Walker, *Note on a Householder implementation of the GMRES method*, Tech. Report UCID-20899, Lawrence Livermore National Lab., CA., 1986.

[27] M. Hoemmen, *Communication-avoiding Krylov subspace methods*, PhD thesis, EECS Dept., U.C. Berkeley, 2010.
[28] D. Imberti and J. Erhel, *Vary the s in Your s-step GMRES*. working paper or preprint, July 2016, `https://hal.inria.fr/hal-01299652`.
[29] W. Joubert and G. Carey, *Parallelizable restarted iterative methods for nonsymmetric linear systems. Part I: theory*, Int. J. Comput. Math., 44 (1992), pp. 243–267.
[30] W. Karush, *An iterative method for finding characteristic vectors of a symmetric matrix*, Pacific J. Math, 1 (1951), pp. 233–248.
[31] S. Kim and A. Chronopoulos, *A class of Lanczos-like algorithms implemented on parallel computers*, Parallel Comput., 17 (1991), pp. 763–778.
[32] S. Kim and A. Chronopoulos, *An efficient Arnoldi method implemented on parallel computers.*, in International Conference on Parallel Processing, vol. 3, 1991, pp. 167–170.
[33] S. Kim and A. Chronopoulos, *An efficient nonsymmetric Lanczos method on parallel vector computers*, J. Comput. Appl. Math., 42 (1992), pp. 357–374.
[34] M. Mehri Dehnavi, Y. El-Kurdi, J. Demmel, and D. Giannacopoulos, *Communication-avoiding Krylov techniques on graphic processing units*, Magnetics, IEEE Transactions on, 49 (2013), pp. 1749–1752.
[35] M. Mohiyuddin, M. Hoemmen, J. Demmel, and K. Yelick, *Minimizing communication in sparse matrix solvers*, in Proc. ACM/IEEE Conference on Supercomputing, ACM, 2009, pp. 36:1–36:12.
[36] V. Simoncini and D. Szyld, *Theory of inexact Krylov subspace methods and applications to scientific computing*, SIAM J. Sci. Comput., 25 (2003), pp. 454–477.
[37] G. Sleijpen, H. van der Vorst, and D. Fokkema, *BiCGstab (l) and other hybrid Bi-CG methods*, Numer. Algorithms, 7 (1994), pp. 75–109.
[38] G. Sleijpen, H. van der Vorst, and J. Modersitzki, *Differences in the effects of rounding errors in Krylov solvers for symmetric indefinite linear systems*, SIAM J. Matrix Anal. Appl., 22 (2001), pp. 726–751.
[39] J. van den Eshof and G. Sleijpen, *Inexact Krylov subspace methods for linear systems*, SIAM J. Matrix Anal. Appl., 26 (2004), pp. 125–153.
[40] H. Van der Vorst and Q. Ye, *Residual replacement strategies for Krylov subspace iterative methods for the convergence of true residuals*, SIAM J. Sci. Comput., 22 (1999), pp. 835–852.
[41] J. Van Rosendale, *Minimizing inner product data dependencies in conjugate gradient iteration*, in Proc. IEEE Internat. Confer. Parallel Processing, IEEE, 1983, pp. 44–46.
[42] H. Walker, *Implementation of the GMRES method using Householder transformations*, SIAM J. Sci. Stat. Comput., 9 (1988), pp. 152–163.
[43] S. Williams, M. Lijewski, A. Almgren, B. Van Straalen, E. Carson, N. Knight, and J. Demmel, *s-step Krylov subspace methods as bottom solvers for geometric multigrid*, in International Symposium on Parallel and Distributed Processing, IEEE, 2014.
[44] I. Yamazaki, H. Anzt, S. Tomov, M. Hoemmen, and J. Dongarra, *Improving the performance of CA-GMRES on multicores with multiple GPUs*, in Proc. IEEE 28th Internat. Parallel and Distributed Processing Symposium, IEEE, 2014, pp. 382–391.
[45] I. Yamazaki, S. Rajamanickam, E. Boman, M. Hoemmen, M. Heroux, and S. Tomov, *Domain decomposition preconditioners for communication-avoiding Krylov methods on a hybrid CPU/GPU cluster*, in Proc. ACM/IEEE Conference on Supercomputing, IEEE, 2014, pp. 933–944.
[46] I. Yamazaki and K. Wu, *A communication-avoiding thick-restart Lanczos method on a distributed-memory system*, in Euro-Par 2011: Parallel Processing Workshops, Springer, 2012, pp. 345–354.